ESTIMATION OF UNCERTAINTY IN CONSTRAINED SPHERICAL DECONVOLUTION FIBER ORIENTATIONS

Ben Jeurissen[1]*, Alexander Leemans*[2]*, Jacques-Donald Tournier*[3,4]*, Jan Sijbers*[1]

[1]Visionlab, Dept. of Physics, University of Antwerp, Belgium
[2]CUBRIC, School of Psychology, Cardiff University, United Kingdom
[3]Brain Research Institute, Melbourne, Australia
[4]Dept. of Medicine, University of Melbourne, Australia

ABSTRACT

Constrained spherical deconvolution (CSD) is a new reconstruction technique that extracts white matter fiber orientations from diffusion weighted MRI data of the brain. However, since these orientations are estimated from noisy data, they are subject to errors, which propagate during fiber tractography. Therefore, it is important to estimate the uncertainty associated with the fiber orientations. In this work, we investigate the performance of a statistical method called the bootstrap, when estimating confidence intervals for CSD fiber orientations. The bootstrap is a nonparametric statistical technique based on data resampling. We used Monte Carlo simulations to measure both its accuracy and precision when applied to CSD. Also, we evaluated an alternative method called the bootknife, which aims to increase the precision of the bootstrap.

Index Terms— Diffusion Weighted MRI; Constrained Spherical Deconvolution; Bootstrap; Confidence Intervals, Monte Carlo simulations

1. INTRODUCTION

Diffusion weighted (DW) magnetic resonance (MR) images of the brain contain information about the orientation of white matter fibers that can be used to study brain connectivity using fiber tractography. Currently, the diffusion tensor model is widely used to extract fiber directions from these data, but fails in regions containing multiple fiber orientations. The constrained spherical deconvolution (CSD) technique has recently been proposed to address this limitation [1]. CSD estimates fiber orientations within each voxel directly from the DW data, using the concept of spherical deconvolution. However, since the fiber orientations are estimated from noisy DW images, they are subject to errors, which propagate in tractography [2]. It has already been shown that the bootstrap method is a very powerful method for characterizing uncertainty in estimates of diffusion tensor imaging (DTI) fiber orientation [3, 4] and it has been successfully used to perform probabilistic DTI tractography [5]. However, this technique has not yet been assessed for CSD. In this work, Monte Carlo simulations will be used to investigate the performance of bootstrap methods in terms of accuracy and precision, when estimating confidence intervals (CI) of the CSD fiber orientations.

2. THEORY & METHODS

2.1. Signal modeling using spherical harmonics

The DW MRI signal is acquired in a set of n_s gradient directions $\{(\theta, \phi)\}$. This signal can be expressed as a linear combination of the real spherical harmonics (SH) $Y_l^m(\theta, \phi)$ of degree l and order m:

$$S(\theta, \phi) = \sum_{l=0}^{L} \sum_{m=-l}^{m=l} c_l^m Y_l^m(\theta, \phi) \quad , \tag{1}$$

where $\{c_l^m\}$ denote the harmonic series coefficients, and L is the maximum harmonic degree. Because only even-degree SH's define symmetric functions, odd-degree harmonics are not included in the representation. Eq.(1) can be expressed as a linear system:

$$\boldsymbol{s} = \boldsymbol{B}\boldsymbol{c} + \epsilon \quad , \tag{2}$$

where $\boldsymbol{B}$ is the $n_s \times n_c$ matrix constructed with the real symmetric SH basis, $\boldsymbol{c}$ is the $n_c \times 1$ vector of even-degree SH coefficients, $\boldsymbol{s}$ is the $n_s \times 1$ DW signal vector and $\boldsymbol{\epsilon}$ is the noise vector. Since only even degrees are used $n_c = (L+1) \times (L+2)/2$. The coefficients $\boldsymbol{c}$ can then be estimated using least-squares minimization:

$$\hat{\boldsymbol{c}} = \boldsymbol{B}^{-1}\boldsymbol{s} \quad . \tag{3}$$

2.2. Fiber orientation estimation using CSD

2.2.1. Spherical deconvolution

From the SH coefficients of the DW signal, the SH coefficients of the fiber orientation distribution (FOD) can be calculated using a technique called spherical deconvolution [6]. This method estimates in each voxel the FOD $F(\theta, \phi)$ by assuming a response function $R(\theta, \phi, \varphi)$ and deconvolving this function from the DW signal $S(\theta, \phi)$. The response function describes the DW signal intensity that would be measured as a function of orientation for a single fiber bundle aligned along the z-axis and can be estimated from the data. The FOD contains information regarding the orientations of the various fiber populations that may be present and their respective volume fractions.

In the SH framework, the spherical deconvolution operation can be performed as:

$$\boldsymbol{f} = \boldsymbol{R}^{-1} \cdot \boldsymbol{c} \quad , \tag{4}$$

where $\boldsymbol{f}$ and $\boldsymbol{c}$ are the $n_c \times 1$ SH coefficient vectors of $F(\theta, \phi)$ and $S(\theta, \phi)$, respectively; $\boldsymbol{R}$ is the $n_c \times n_c$ rotational harmonic matrix of $R(\theta, \phi, \psi)$. Similarly to the spherical harmonics, the rotational harmonics form a complete orthonormal basis over the space of pure rotations. Due to the rotational symmetry of $R(\theta, \phi, \psi)$, the matrix $\boldsymbol{R}$ can be shown to be diagonal. Details on its calculation can be found in [1].

The deconvolution operation, however, is sensitive to noise. Until recently, a low-pass filter was used, attenuating the high frequency coefficients of $S(\theta, \phi)$. This removes high angular frequency

noise, but also reduces angular resolution of the reconstructed FOD. Fig.1(a) shows the FOD of two fiber populations intersecting at an angle of 60° reconstructed using filtered spherical deconvolution.

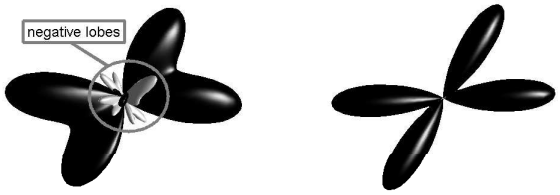

(a) Filtered (b) Constrained

Fig. 1. Filtered vs. constrained spherical deconvolution.

2.2.2. Constrained spherical deconvolution

Recently, a new technique, called CSD was proposed, which does not require low-pass filtering [1]. This method follows from the observation that standard spherical deconvolution results in spurious negative lobes in the FOD (Fig.1(a)), which are physically impossible. By introducing a constraint that minimizes these negative lobes, the effects of noise can be reduced without a low-pass filter while preserving angular resolution (Fig.1(b)).

In brief, the method involves the following steps. First, an initial estimate of the FOD is obtained using filtered spherical deconvolution. A set of directions is then identified, along which the FOD amplitude is negative. This information is then incorporated as a Tikhonov constraint, driving the amplitude of the FOD along those orientations to zero. An improved estimate of the FOD is then obtained by solving the Tikhonov problem, providing a new set of negative amplitude directions. The procedure is repeated until convergence is achieved. A more detailed explanation of these steps can be found in [1].

In this work, we will use CSD to extract the FOD from the DW signal in each voxel.

2.3. Bootstrap methods for estimating uncertainty

If CSD is performed on repeated acquisitions of the same fiber orientations, this will result in a different FOD for every acquisition as a result of different instances of noise. To estimate this variability we could measure a very large number of acquisitions and calculate for example a 95% CI. However, in practice only a few data sets can be acquired and this technique would yield very poor results. Therefore, in order to estimate the uncertainty of CSD fiber orientations we have to resort to more advanced techniques, as described in the following sections.

2.3.1. Bootstrap

Recent studies have shown the potential of the bootstrap method to estimate the reproducibility of in vivo white matter orientations derived from diffusion tensor MRI [4]. The bootstrap [7] is an empirical, nonparametric, statistical technique based on data resampling. Bootstrap replaces complicated and often inaccurate approximations to uncertainty measures, e.g. bias and variance, with computer simulations based on real data.

The method requires the acquisition of N repeats of a complete DW data set, so that N samples are available for each gradient direction. A resampled DW data set can then be produced by randomly selecting N samples with replacement for each direction. Consider N repeated data sets of U DW measurements. Each complete DW acquisition is a $U \times 1$ vector, and we repeat the acquisition N times. The total data available to us is a matrix D, of dimensions $U \times N$, where each column in D is one complete acquisition. To generate a single bootstrap realization, we create a new matrix D_b, of dimension $U \times N$, where each row contains N values randomly sampled with replacement from the corresponding row in D (Fig.2(a)). In this way, a full $U \times N$ DW data set is produced from a random combination of the images in the N repeats of the original data set, which can then be processed using the method under investigation, in our case CSD. By repeating this procedure N_b times, we obtain N_b estimates of the FOD and the associated peak orientations, which can be used to estimate the reproducibility of the reconstructed fiber orientations (Fig.2(b)).

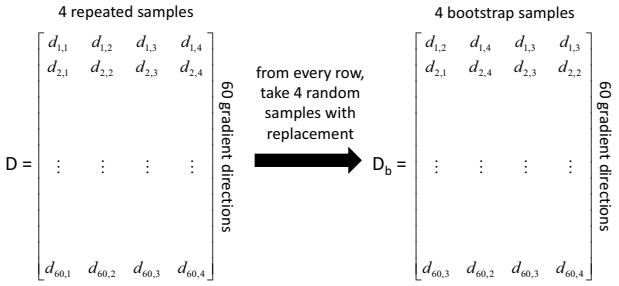

(a) One bootstrap realization.

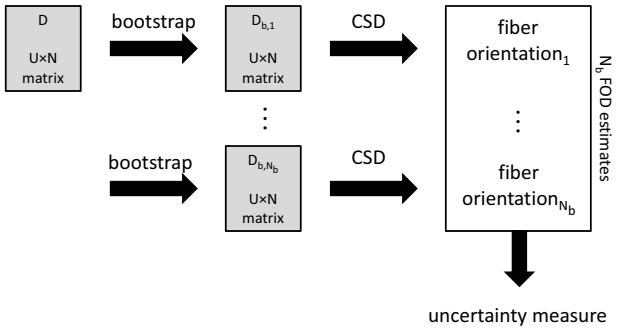

(b) Estimate uncertainty from many bootstrap realizations.

Fig. 2. Example of the bootstrapping procedure with $U = 60$ gradient directions and $N = 4$ repeated acquisitions.

2.3.2. Bootknife

When the number of repeats N is small, bootstrap-estimated uncertainties are noticeably downwardly biased, in the same way that the uncorrected variance is not an unbiased estimator of the true population variance. To remedy this bias, the bootknife method was proposed [8]. Basically, the bootknife is a combination of the jackknife and bootstrap. Prior to selecting one of the N available samples for each direction, we eliminate one measurement at random from each row in D, giving us a matrix D_j of dimension $U \times (N-1)$ (jackknife). We then create a bootstrap matrix D_b of dimension $U \times N$,

but because we are choosing N samples from a row of length $N-1$, we guarantee that at least one of the measurements will be repeated. This technique has already been used successfully on DTI [9].

3. EXPERIMENTS

Simulation experiments were performed to compare the CSD bootstrap and bootknife estimate of fiber orientation uncertainty to a gold standard.

Two diffusion tensor profiles at angles ranging from 60° to 90° were combined to simulate the noiseless DW signal for a two fiber population:

$$S(\boldsymbol{u}) = S_0 e^{-b\boldsymbol{u}\boldsymbol{D}_1\boldsymbol{u}^T} + S_0 e^{-b\boldsymbol{u}\boldsymbol{D}_2\boldsymbol{u}^T} \quad , \tag{5}$$

where S_0, the non-DW signal, was set to 1 without loss of generality. The diffusion weighting b was set to $3000s/mm^2$. Sixty diffusion-encoding gradient directions $\boldsymbol{u}$ were used, distributed evenly on the half sphere [10]. This setup corresponds to a realistic high angular resolution DW acquisition. Both diffusion tensors $\boldsymbol{D}_i (i = 1, 2)$ had a fractional anisotropy (FA) of 0.8. The mean apparent diffusion coefficient (ADC) was set to $600 \times 10^{-6} mm^2/s$. Rician noise was added to give a SNR (of S_0) of 15 to 40, which is the clinical range. This experiment was repeated 10,000 times. The FOD was calculated for every DW signal, using CSD with harmonic degree $L = 8$. From these FOD's, peaks were extracted using a quasi-Newton optimization method. The average peak directions were calculated as the first eigenvector of the mean dyadic tensor of all 10,000 peak directions [3]. Finally, the 95% CI of the angular deviation between the individual and average peak orientations was calculated, representing the "cone of uncertainty" [4] around the average peak orientation. Fig.3 shows a visualization of such cones for different SNR values. Note that the cones are wider at low SNR, indicating a higher uncertainty of the fiber orientations.

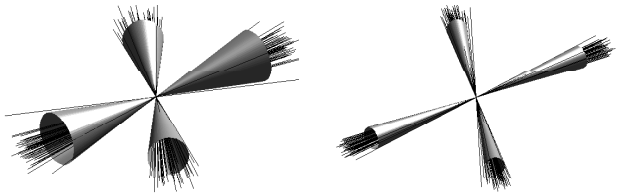

(a) SNR=15 (b) SNR=30

Fig. 3. Example of a 95% CI of two fiber orientations with inter-fiber angle of 60° at different SNR levels. For clarity, the number of repeats was only 60. The cones represent the 95% CI, the black lines the fiber orientation estimates from the individual repeats. This means 95% of the black lines lie within the cone. N was set to 1.

3.1. Bootstrap

Nine bootstrap experiments were considered, with the number of repeated acquisitions, N, ranging from 2 to 10. For each bootstrap design, we derived 1000 bootstrap realizations of the FOD. Fiber orientations were extracted as described above. To determine the effect of the number of bootstrap realizations on the estimated fiber orientations, N_b was incremented from 100 to 1000 in steps of 100. The entire procedure was repeated 100 times to determine the precision of a particular bootstrap experiment. Mean and standard deviation of the 95% CI (across the 100 repeats) were computed.

3.2. Bootknife

The experiments for the bootknife were performed in the same way as described in subsection 3.1.

4. RESULTS

Fig.4 shows the bias of the 95% CI relative to the gold standard 95% CI as a function of the number of repeated measurements N. The number of bootstrap realizations N_b was fixed to 1000 and SNR was set to 25. The dotted lines represent the bias of the bootstrap estimates. The dashed lines show the bias of the bootknife estimates. Only inter-fiber angles of 60° (Fig.4(a)) and 90° (Fig.4(b)) are shown. Note that other inter-fiber angles and SNR levels yielded similar results, but were left out for clarity. As expected, the CI's are significantly underestimated by the bootstrap when the number of repeated acquisitions N is small. Accuracy can be improved by increasing N, but there still remains a negative bias even at $N = 10$ repeated measurements. The bootknife estimates on the other hand tend to be very close to the gold standard over the entire range of N and at different inter-fiber angles.

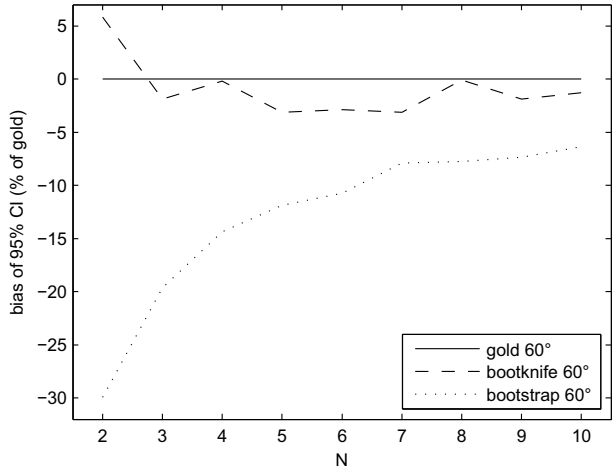

(a) Two fibers at 60°

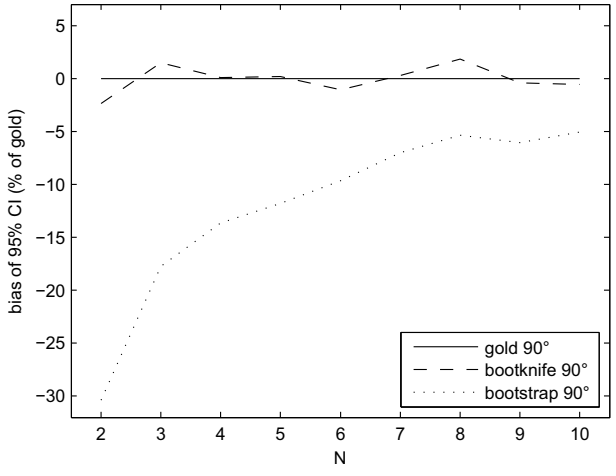

(b) Two fibers at 90°

Fig. 4. Relative bias of mean bootstrap 95% CI as a function of N.

Fig.5 shows the relative standard deviation over 100 bootstrap experiments of the 95% CI's as a function of the number of bootstrap realizations N_b. The number of repeated experiments N was fixed to 6 and SNR was set to 25. The dotted lines represent the bootstrap estimates. The dashed lines show the bootknife estimates. Smaller standard deviations are better. Only inter-fiber angles of 60° (indicated by $\times$) and 90° are shown. Note that other inter-fiber angles and SNR levels yielded similar results, but were left out for clarity. The plot shows improvement in precision of the CI's by increasing the number of bootstrap realizations N_b, but increasing the number of bootstrap iterations only seems sensible up to approximately 700.

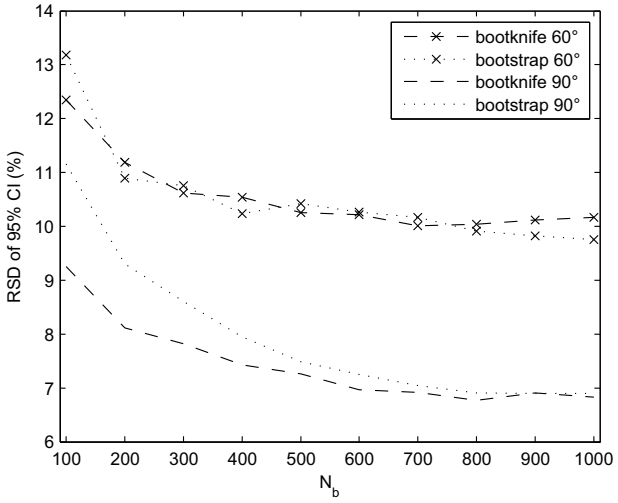

Fig. 5. Relative standard deviation (RSD) of bootstrap 95% CI as a function of N_b.

5. DISCUSSION & CONCLUSION

In this work, we investigated the performance of the bootstrap method in terms of accuracy and precision when estimating confidence intervals of CSD fiber orientations and compared it to an alternative bootstrap method, called bootknife.

The precision of the bootstrap and bootknife method depends on the number of bootstrap realizations and our results show that the number of bootstrap iterations should be 700, which is not a problem since it doesn't impact acquisition time. On the other hand, the accuracy of the bootstrap and bootknife method depends on the number of repeated acquisitions. Our results show that the "classic" bootstrap significantly underestimates the uncertainty when few repeated acquisitions are available. This is in accordance with earlier studies that were performed using DTI [11, 9]. However, high angular resolution diffusion imaging data, like CSD data, typically have very few repeated acquisitions available. While it may be tempting to use this bootstrap procedure with just a few repeats, our results show that this yields poor accuracy. Using the bootstrap for probabilistic tractography can thus have considerable consequences, since the error introduced by this bias will produce tracts that do not represent the true variability inherent in the data.

We also showed that the downward bias of the "classic" bootstrap for CSD can be removed using the bootknife approach. This allows good CI estimates and probabilistic tractography, using only a few repeated acquisitions. However, in a clinical setting, even two repeated measurements can already render acquisition time unacceptably long. Our future research will therefore focus on model-based bootstrapping methods. These methods fit the DW signal to a model and perform bootstrapping on the residuals, generating an arbitrary amount of bootstrap realizations from just one acquisition, drastically reducing the acquisition time. Such methods have already been used successfully to estimate uncertainty of DTI fiber orientations [9, 12] and very recently, Q-Ball Imaging (QBI) fiber orientations [13].

6. REFERENCES

[1] J.D. Tournier, F. Calamante, and A. Connelly, "Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution," *NeuroImage*, vol. 35, pp. 1459–1472, May 2007.

[2] T.E. Behrens, M.W. Woolrich, M. Jenkinson, H. Johansen-Berg, R.G. Nunes, S. Clare, P.M. Matthews, J.M. Brady, and S.M. Smith, "Characterization and propagation of uncertainty in diffusion-weighted MR imaging," *Magn Reson Med*, vol. 50, pp. 1077–1088, November 2003.

[3] S. Pajevic and P.J. Basser, "Parametric and non-parametric statistical analysis of DT-MRI data," *J. Magn. Reson.*, vol. 161, pp. 1–14, March 2003.

[4] D.K. Jones, "Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI," *Magn Reson Med*, vol. 49, pp. 7–12, January 2003.

[5] D.K. Jones and C. Pierpaoli, "Confidence mapping in diffusion tensor magnetic resonance imaging tractography using a bootstrap approach," *Magn Reson Med*, vol. 53, pp. 1143–1149, May 2005.

[6] J.D. Tournier, F. Calamante, D.G. Gadian, and A. Connelly, "Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution," *NeuroImage*, vol. 23, pp. 1176–1185, November 2004.

[7] B. Efron, "Bootstrap Methods: Another Look at the Jackknife," *Annals of Statistics*, vol. 7, no. 1, pp. 1–26, 1979.

[8] T. C. Hesterberg, "Unbiasing the bootstrap-bootknife sampling vs. smoothing," in *Proceedings of the American Statistical Association*, 2004, pp. 2924–2930.

[9] S. Chung, Y. Lu, and R.G. Henry, "Comparison of bootstrap approaches for estimation of uncertainties of DTI parameters," *NeuroImage*, vol. 33, pp. 531–541, November 2006.

[10] D.K. Jones, M.A. Horsfield, and A. Simmons, "Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging," *Magn Reson Med*, vol. 42, pp. 515–525, September 1999.

[11] R.L. O'Gorman and D.K. Jones, "Just how much data need to be collected for reliable bootstrap DT-MRI?," *Magn Reson Med*, vol. 56, pp. 884–890, October 2006.

[12] B. Whitcher, D.S. Tuch, J.J. Wisco, A.G. Sorensen, and L. Wang, "Using the wild bootstrap to quantify uncertainty in diffusion tensor imaging," *Hum Brain Mapp*, Apr 2007.

[13] J.I. Berman, S. Chung, P. Mukherjee, C.P. Hess, E.T. Han, and R.G. Henry, "Probabilistic streamline q-ball tractography using the residual bootstrap," *NeuroImage*, vol. 39, pp. 215–222, January 2008.